

# HPCBS

## High Performance Commercial Building Systems

### Advanced Nonintrusive Monitoring of Electric Loads

*Element 5 - Integrated Commissioning and Diagnostics*  
*Project 2.2 - Monitoring and Commissioning of Existing Buildings*

Massachusetts Institute of Technology

IEEE Power and Energy  
Magazine, March/April, 2003. 56-63.



## **Acknowledgement**

This work was supported by the California Energy Commission, Public Interest Energy Research Program, under Contract No. 400-99-012 and by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

### **DISCLAIMER**

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

This report was prepared as a result of work sponsored by the California Energy Commission (Commission). It does not necessarily represent the views of the Commission, its employees, or the State of California. The Commission, the State of California, its employees, contractors, and subcontractors make no warranty, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the use of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the Commission nor has the Commission passed upon the accuracy or adequacy of the information in this report.

# Advanced Nonintrusive Monitoring of Electric Loads

Christopher Laughman    Douglas Lee    Robert Cox  
Steven Shaw    Steven Leeb    Les Norford    Peter Armstrong

## I. INTRODUCTION

**T**HE boom in communications research and development has made highly sophisticated means for gathering, moving, and exchanging information part of everyday life. Examples include the Internet, Bluetooth and 802.11b networks, cellular, and satellite communications. However, the process of gathering and analyzing the data to send over this expanding web of information networks remains expensive for many applications. A building facilities manager may be able to configure a cellular phone to monitor the energy usage and operating schedule of an HVAC plant, but the usefulness of the information transmitted will generally be directly proportional to the complexity and size of the installed sensor array.

This difference between information gathering and information networking serves to re-emphasize an established fact: though remote access to information and control inputs may be obtained easily and inexpensively via networking, access does not provide useful information without installation of a potentially expensive and intrusive sensor array. Mass production may ultimately reduce sensor cost, especially for solid-state or technologically advanced micro-electromechanical sensors. Installation expenses are likely to remain high, especially for temporary monitoring applications where data is gathered for a brief window of time with a removable sensor network. In addition, the reliability of a monitoring system with many sensors may be reduced in comparison to a system with relatively fewer sensors.

The Nonintrusive Load Monitor (NILM) can determine the operating schedule of electrical loads in a target system from measurements made at a centralized location, such as the electric utility service entry. In contrast to other systems, the NILM reduces sensor cost by using relatively few sensors. The NILM disaggregates and reports the operation of individual electrical loads like lights and motors using only measurements of the voltage and aggregate current at the utility service point of a building. It can also identify the operation of electromechanical devices in other kinds of power distribution networks. For example, the NILM can determine the load schedule in an aircraft from measurements made only at the generator, or in an automobile from measurements made at the alternator/battery block. The NILM can distinguish loads even when many are operating at one time.

The NILM is an ideal platform for extracting useful information about any system that uses electromechanical devices. The NILM has a low installation cost and high reliability because it uses a bare minimum of sensors. It is possible to use modern state and parameter estimation algorithms to verify remotely the “health” of electromechanical loads by using a NILM to associate measured waveforms with the operation of individual loads. The NILM can also monitor the operation of the electrical distribution system itself, identifying situations where two or more otherwise healthy loads interfere with each other’s operation through voltage waveform distortion or power quality problems.

Strategies for nonintrusive monitoring have developed over the last twenty years. Advances in computing technology make a new wealth of computational tools useful in practical, field-based NILM systems. This paper reviews techniques for high performance nonintrusive load and diagnostic monitoring and illustrates key points with results from field tests.

## II. BACKGROUND AND EARLY APPROACHES

One of the earliest approaches to nonintrusive monitoring, developed in the 1980’s at MIT by Professor Fred Schwegge and Dr. George Hart, had its origins in load monitoring for residential buildings. Under Dr. Hart’s ingenious scheme, the operating schedules of individual loads or groups of loads are determined by identifying times at which electrical power measurements change from one nearly constant (steady-state) value to another. These steady-state changes, known as events, nominally correspond to the load either turning on or turning off, and are characterized

by their magnitude and sign in real and reactive power. Recorded events with equal magnitudes and opposite signs are paired to establish the operating cycles and energy consumption of individual appliances. This process of detecting steady-state changes provides the basis for a commercial version of this early work, as described in a 1999 article in the IEEE Computer Applications in Power.

The 1999 CAP article [2] describes a five-step process for load disaggregation through the detection of changes in aggregate power consumption. First, an edge detector is used to identify changes in steady-state levels. Second, a cluster analysis algorithm is used to locate these changes in a two-dimensional “signature space” of real and reactive power (a  $\Delta P$ - $\Delta Q$  plane). The signature space reduces the potentially complicated load transient data to a two-dimensional space of changes in power consumption with a pleasing and useful graphical interpretation. Third, positive and negative clusters of similar magnitude are paired or matched (especially for “two-state” loads that turn on, consume a relatively fixed power, and turn off). In a fourth step known as anomaly resolution, unmatched clusters and events are paired or associated with existing or new clusters according to a best likelihood algorithm. In the fifth and final step, pairs of clusters are associated with known load power consumption levels to determine the operating schedule of individual loads. This step uses information gathered during a training or survey phase in the building.

This five-step approach to nonintrusive load monitoring has the advantage of making intuitive sense, and has been demonstrably successful in certain classes of buildings, such as residences. Recent field tests in a variety of other buildings have uncovered a number of limitations of the technique of examining steady-state changes in a signature space. Some of these limitations are well established, while others are relatively new.

The two-dimensional signature-space technique relies on at least three key assumptions that limit its effectiveness. The first assumption is that different loads of interest exhibit unique signatures in the  $\Delta P$ - $\Delta Q$  plane. They may not, especially in commercial and industrial facilities where the number and variety of loads is generally greater than in a residence. The two-dimensional signature-space becomes crowded with indistinguishable loads as the number and kind of loads increases. In a home there may be only one 5 kW load, such as a hot water heater. In comparison, a light commercial facility with an HVAC plant, office equipment, and lighting may have several different loads that overlap ambiguously in the  $\Delta P$ - $\Delta Q$  plane.

The second limiting assumption is that load composition is determined by steady-state power consumption. This requires waiting until transient behavior settles out so that steady state values can be measured. We have found it difficult to find any suitable time scale in industrial or commercial environments that yields reliable steady-state measurements. A conservative estimate of steady-state power might require a long interval of near-constant demand. Requiring a long steady-state waiting time prevents the monitor from tracking rapid sequences of loads. Some loads will be missed, and these may not be caught in the anomaly resolution phase. Short waiting times, on the other hand, may trigger measurements in the middle of load transients, resulting in spurious events in the cluster analysis phase. Field studies have shown that large HVAC loads such as fans and chillers might take from 30 seconds to several minutes to gradually spin-up to their final operating speed. Other loads, such as variable speed drives (VSDs), may never settle to a steady-state. Furthermore, if the variations in power consumption are large enough, VSDs and similar loads could prevent the monitor from finding a steady-state consumption level or recording any events. Some industrial loads that ordinarily settle to steady state conditions may fail to do so if they include poorly tuned controllers or other faults.

A third limitation is that most steady-state NILM systems process data in “batch” format using a day or more of stored data. This is not strictly implied by the five-step disaggregation procedure outlined above, but is based on the assumption that near-real-time identification of load operation is not necessary. This limits the monitor to load survey and power scorekeeping applications, excluding the vast potential for applications in real-time diagnostics.

### III. ADVANCED TECHNIQUES FOR MONITORING

Steady-state monitoring techniques are successful in homes and small businesses because of the low event generation rate and number of loads at these sites. Medium to large size commercial and industrial facilities require a more sophisticated approach, due in part to high rates of event generation, load balancing, and power factor correction. Savvy commercial and industrial facilities managers want near-real-time monitoring information in addition to batch results over an interval of several weeks. Real-time identification is also essential if the NILM is to serve as a “black-box” diagnostic monitor in transportation systems like aircraft or naval vessels. Field tests have demonstrated the value of a number of enhancements that enable an advanced NILM to tackle complex monitoring environments.




Fig. 1. Plot of the steady state power consumption of a computer and a bank of incandescent lights. Real power is plotted against reactive power.

### A. Higher Harmonics

Higher harmonics in the aggregate current signal can be used to distinguish loads with overlapping clusters in the  $\Delta P$ - $\Delta Q$  signature space. Many loads draw distorted, non-sinusoidal current due to their inherent physical characteristics or the presence of power electronics. Examples include office equipment (i.e., computers and copiers), actively controlled industrial equipment such as variable speed fans, and electroplating baths. Our advanced field monitoring system uses a phase-locked short-time Fourier transform of current waveforms collected at sample rates of 8000 Hz or higher to compute “spectral envelopes” that summarize time-varying harmonic content. For a single phase load, real and reactive power correspond to the envelopes of in-phase and quadrature current drawn by the load relative to the voltage. The short-time Fourier transform computes estimates of the real, reactive and higher frequency components of the current.

Figure 1 shows that higher harmonics are useful for distinguishing loads that are similar in the P-Q signature space. Figure 1 shows a collection of turn-on and turn-off events recorded at a site with personal computers and a bank of incandescent lamps. The loads are practically indistinguishable in the P-Q signature space because they consume essentially the same real and reactive power. However, typical computer power supplies draw “signature” third harmonic currents. The loads are easily separable by examining higher harmonics.

Figures 2 and 3 illustrate the relative ease of distinguishing individual loads in a three-dimensional space with axes denoting changes in real power, reactive power, and third harmonic. Our advanced NILM routinely examines harmonic content up to and including eighth harmonic, and can be customized to examine higher harmonics as necessary.

### B. Transient Detection

Our advanced NILM recognizes individual loads based on distinctive load transient shapes. The transient behavior of a typical load is intimately related to the physical task that the load performs. For example, the turn-on transients associated with a personal computer and with an incandescent lamp are distinct because charging capacitors in the computer power supply is fundamentally different from heating a lamp filament. Overall transient profiles tend to be preserved even in loads that use active waveshaping or power factor correction. Most loads observed in the field have repeatable transient profiles, or at least sections of the transient profile that are repeatable. Transient-based recognition permits near-real-time identification of load operation, especially turn-on events.

Transients are identified by matching events in the incoming aggregate power stream to previously defined transient signatures, or “exemplars.” Exemplars can be determined, for example, by a one-time direct observation of the device in question, or by previous training in the laboratory. Pre-training has proven to be a reasonable approach for very repeatable loads that show up in large quantities, such as fluorescent lamp ballasts. The exemplar may be composed of multiple parts for transients with a number of distinct sections. Each section of the exemplar can be shifted in time and offset to match an incoming transient data. In addition, an overall gain may be applied to all sections of the exemplar to




Fig. 2. Computer and incandescent lightbulb turn-on and turn-off event clusters in the  $\Delta P$ ,  $\Delta 3^{rd}$  harmonic,  $\Delta Q$  coordinate system.




Fig. 3. Three cross-sections of the  $\Delta P$ ,  $\Delta 3^{rd}$  harmonic,  $\Delta Q$  coordinate system, illustrating the separation of clusters.

achieve a better fit. Each event detected is compared to the full set of exemplars by using a least squares criterion to select the appropriate shifts and gains. The match with the lowest residual norm per number of points is then compared to a threshold. If the fit is good enough, the event is classified as a match to the exemplar. If not, the event is left unclassified. Correct classification of overlapping transients is possible using properly designed exemplars.

The value of examining transient information can be understood graphically by examining Figures 4, 5, and 6. Figures 4 and 5 show the spectral envelopes corresponding to the fundamental, in-phase component of current (essentially, “real power”) demanded in four different buildings in California. The top graph in Figure 4 shows the power consumed solely by the HVAC panel in a building (“ISD”) in Los Angeles County. The bottom trace of Figure 4 and the top trace of Figure 5 show power consumption in two different public schools near San Francisco (“Hanna” and “Pinole”, respectively). Finally, the bottom trace in Figure 5 shows the total power consumption in another building in Los Angeles County (“Comm”); the inset in this plot is a magnified portion of the same waveform, illustrating the high density of events. These four traces show a progression in increasing event generation in a building HVAC panel using steady-state signature detection. As the complexity and variation on the electrical grid increases, it becomes more difficult to tune a steady-state change detector to function at all. For buildings with behavior nearing that of the Comm building, it becomes essential to either augment the steady-state approach or discard it completely in favor of transient identification. By enabling the monitor to deal with higher rates of event generation, a single NILM using transient event detection can generally monitor a comparatively more complex building power network than with steady-state change detection alone.

In addition to better load disaggregation, transient analysis can provide diagnostic information. Diagnostics can be developed by exploring the relationship between the electrical transient and the physics of the load. Specifically, a load model can often be developed which relates the shape of the electrical transient to physical or design parameters of the load. The load model is often a differential equation, although we have experimented with other choices. It is often possible to deduce the physical condition of the load by examining load model parameters when the transient data is rich enough to identify the load model.

Figure 6 illustrates the diagnostic capabilities of our advanced NILM. The numbers displayed in the text window below the graphs in Fig. 6 show estimated parameters of an induction machine model. The solid lines in the graph represent raw data, while the points show the results of a simulation with the estimated parameters. The close agreement shown in the figure indicates that the parameters and model accurately predict the response of the actual load. In addition to developing these physically-based models, we have also developed techniques that can accurately estimate model parameters **without** a good initial guess. These estimation methods enable a building manager to trend load parameters and use parameter information to predict impending faults. In the case of an induction motor, such faults might include shorted motor windings, broken rotor bars, and especially failures in a mechanical system attached to the shaft, e.g., a slipping belt.




Fig. 4. Plots of the aggregate energy consumption collected at the HVAC service entry in a California county government building and at the main service entrance of a California elementary school.




Fig. 5. Plots of the aggregate energy consumption collected at the main service entries of a California middle school and a communications service building. The inset plot illustrates the high density of events; this data in this plot is averaged in order to enhance the visibility of the individual events.




Fig. 6. Startup transient for an induction machine.

### C. Disaggregating Continuously Variable Loads

A third benefit obtained through transient identification and spectral analysis is that our advanced NILM has the ability to monitor buildings with continuously variable loads. Many such loads, such as the VSDs used in HVAC fans, servo their electrical power consumption under the influence of an active closed-loop controller. Transient detection can simplify the tracking of these and other continuously variable loads in several ways.

Figure 7 shows three spectral envelopes associated with the operation of a 40 Hp VSD in an HVAC plant on the

MIT campus. The top trace in this figure shows the real power demanded by a variable-speed fan drive in an HVAC system. The drive begins with an “open-loop” spin-up to operating speed during the first 40 seconds of operation. This open-loop spin-up is repeatable because a microprocessor controls the startup profile every time the drive is activated. A NILM can recognize that a VSD is active in the building using transient recognition, but transient recognition does not provide a means for continuously tracking the power consumption of a variable load like a VSD.

When the NILM recognizes a continuously variable load like a VSD, we have found that it is often possible to disaggregate the variable load by carefully examining the spectral envelopes. The VSD connects to the utility through a three-phase, delta-connected rectifier. The three-phase rectifier draws distorted, pulsatile current waveforms from each of the three utility phases. As shown in the lower two traces in Figure 7, this rectifier set has the effect of causing the VSD to create characteristic traces not only in real power but also in the fifth and seventh harmonics, as illustrated in the middle and bottom traces respectively. These higher harmonic traces have envelopes that roughly track the shape of the real power trace. If the VSD is the only load at a target site which generates the bulk of a particular harmonic signal (such as the seventh harmonic) in the aggregate data, this particular harmonic signal can be functionally mapped and subtracted from the “real power” spectral envelope in order to remove the continuously varying component of the “real power” signal due to the VSD from the aggregate power signal. This makes it easier to identify the remaining active loads in the aggregate power signal. In general, the higher harmonic information can be used both to track the energy consumption of the variable load and also to disaggregate this variable load from the power consumption of the remaining loads. There are many possible variations on this technique of refining the mappings between a particular harmonic and other harmonics; these techniques allow one group of loads to be disaggregated from another.




Fig. 7. Startup transient of a variable speed drive.

Additionally, harmonic signatures can provide direct diagnostic information even without parameter estimation. Returning to Figure 7, the drive is operating under closed loop control as it attempts to regulate the pressure in a distant duct by varying fan speed starting at 100 seconds. The control loop is poorly tuned and the drive exhibits a slow, large amplitude oscillation in “steady state.” This pathology is not easily caught in a brief inspection of the drive, as the oscillation is fairly slow; nevertheless, the identification of this failure mode is important because it wastes energy and wears the mechanical components of the drive. This condition is easily caught by an advanced NILM which examines transients and spectral envelopes.



#### IV. APPLICATIONS AND THE FUTURE

We have implemented all of these analytical enhancements for nonintrusive monitoring on an inexpensive computer platform for use in the field. A suite of NILM tools, including sophisticated model-based diagnostic algorithms that track (and could in principle trend) model parameters to determine the health of critical loads, has been developed under the Linux operating system environment. All of these software tools run on a Pentium-class personal computer (300 MHz clock or higher) with a PCI-bus data-acquisition card; any relatively inexpensive personal computer, laptop, or embedded system like a PC104 chassis could be used to develop a modern monitor which performs transient event detection. We have used all of these platforms with success in buildings and transportation systems.

Current experiments in the field are aimed at enhancing the diagnostic capabilities of the NILM, integrating it into building energy management and control systems, and improving it as a load monitor for conducting surveys and energy-usage scorekeeping. Even in very modern buildings with sophisticated control and consumption metering systems, we have found that the monitor is invaluable to sophisticated facilities managers. As the NILM requires virtually no additional wiring effort and few installed sensors, it provides invaluable building operating history and health information in a package which is both easy to install and highly reliable.

In the future, we expect to introduce new event detection and classification schemes where load disaggregation and usage tracking are increasingly accurate. The best approaches appear to blend all the tools described above to some degree. For example, the highest accuracy in usage disaggregation may be obtained by combining the steady-state and transient approaches to event recognition. Additionally, the steady-state signature scheme can obviously be enhanced by working in a larger signature space with harmonic data. Yet another combination of steady-state and transient identification methods may be used to alleviate the problem of defining steady-state intervals by using transient event detection to tease apart overlapping transients which occur at high rates. Furthermore, the transient detection approach may be minimally successful when transient shapes are masked or distorted by background, periodic oscillations in the building; by examining changes in the steady-state sections of the aggregate data stream, the transient event detector may be more successful in looking for “most likely” shapes in a cluster of overlapping or distorted transients. We are currently developing “arbiters” that use many different approaches to identify events with the highest possible accuracy and precision.

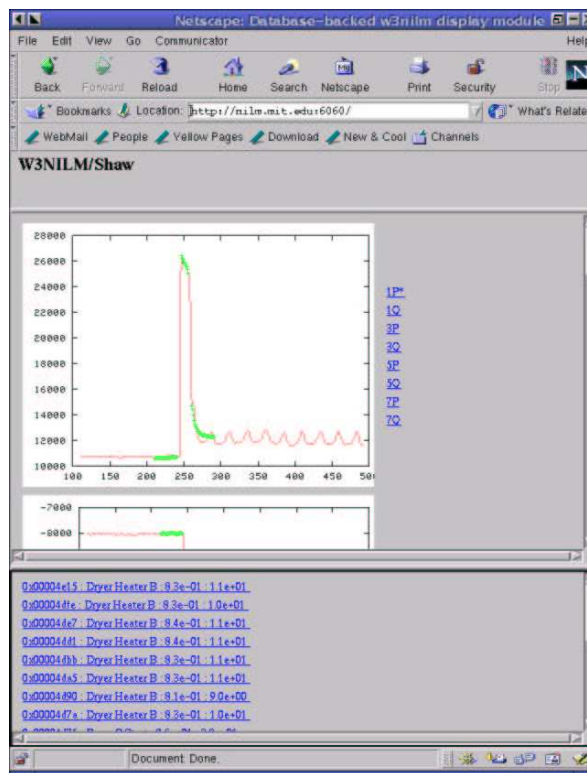


Fig. 8. Nonintrusive monitoring on the web.

Because the advanced NILM uses a conventional computer platform, it is easily connected with communications networks, including the web. An example from a NILM installed on the MIT campus is shown in Figure 8. The Netscape window shows a turn-on transient associated with a washing machine on top of an 11 kW base load; this can be viewed via a remote connection to a computer monitoring a campus dormitory containing a laundromat. The cusp-like steady-state, following the initial acceleration of the agitator motor to operating speed, is due to the “steady-state” oscillation of the agitator. Other load transients recognized by the monitor, including several dryers, are listed in the text window below the graphs. The items in this list are hypertext links corresponding to recorded events; these links can be selected in order to view a record for each indicated transient, including time, energy consumption, and diagnostic parameters (e.g. heater resistance, motor parameters like stator resistance, magnetizing and leakage inductances, etc.) We anticipate that the advanced NILM technology will provide an ideal platform for remote monitoring and for customizing network-accessible information packages for utility customers and building occupants.

#### ACKNOWLEDGMENTS

This research was funded by the California Energy Commission and a grant from the Grainger Foundation.

#### REFERENCES

- [1] “Wired and Wireless,” *Technology Review Magazine*, Volume 104, No. 5, June 2001, pp. 40–82.
- [2] Drenker, S. and A. Kader, “Nonintrusive Monitoring of Electrical Loads,” *IEEE Computer Applications in Power*, Volume 12, Issue 4, pp. 47–51.
- [3] Hart, G.W., “Nonintrusive Appliance Load Monitoring,” *Proceedings of the IEEE*, Volume 80, No. 12, pp. 1870–1891, December 1992.
- [4] Leeb, S.B., S.R. Shaw, J.L. Kirtley, Jr., “Transient Event Detection in Spectral Envelope Estimates for Nonintrusive Load Monitoring,” *IEEE Transactions on Power Delivery*, Volume 10, Number 3, pp. 1200–1210, July 1995.
- [5] Norford, L.K. and S.B. Leeb, “Non-Intrusive Electrical Load Monitoring in Commercial Buildings Based on Steady State and Transient Load-Detection Algorithms,” *Energy and Buildings*, Volume 24, pp. 51–64, 1996.
- [6] Shaw, S.R., C.B. Abler, R.F. Lepard, D. Luo, S.B. Leeb, and L.K. Norford, “Instrumentation for High Performance Nonintrusive Electrical Load Monitoring,” *ASME Journal of Solar Energy Engineering*, Vol. 120, No. 3, August 1998, pp. 224–229.
- [7] Shaw, S.R. and S.B. Leeb, “Identification of Induction Motor Parameters from Transient Stator Current Measurements,” *IEEE Transactions on Industrial Electronics*, Volume 46, No. 1, February 1999, pp.139–149.
- [8] Shaw, S.R., R.F. Lepard, S.B. Leeb, and C.R. Laughman, “A Power Quality Prediction System,” *IEEE Transactions on Industrial Electronics*, Volume 47, No. 3, June 2000, pp. 511–517.

PLACE  
PHOTO  
HERE

**Christopher R. Laughman** received the S.B and the M.Eng degrees from the Massachusetts Institute of Technology in 1999 and 2001, respectively. He is currently a research engineer in the Laboratory for Electromagnetic and Electronic Systems. Current research includes the investigation of modelling and parameter estimation problems in buildings and transportation systems.